

REMARKS

Claims 1-67 are pending in the application, of which Claims 1, 11, 13, 32, 33, 41, 46, 50, 51, 52, 53, 54, 62 and 67 are independent. Claims 1-3, 5, 7-15, 17, 19-27, 29-34, 36, 37, 39-43, 45-47, 49, 50 and 53-67 were rejected. The Applicants note with appreciation that Claims 51 and 52 were allowed, and Claims 4, 6, 16, 18, 28, 35, 38, 44 and 48 were objected as being dependent on a rejected claim, but would be allowable if rewritten in independent form. Claims 1, 11, 13, 29, 32, 33, 41, 45, 46, 50, 53, 54, 62, 66 and 67 are amended by the present amendment. For the reasons described below, all claims are in condition for allowance.

Claim Amendments

Claims 1, 11, 13, 32, 33, 41, 46, 50, 53, 54, 62 and 67 are amended to emphasize that the member nodes are members in the cluster and can thus access the cluster definition on the shareable storage regardless of network connectivity. Support for this amendment can be found throughout the application as originally filed.¹ Further, Claims 1, 11, 13, 29, 32, 33, 41, 45, 46, 50, 53, 54, 62, 66 and 67 are amended to specify that there are at least two member nodes in the cluster. Acceptance is respectfully requested.

Rejections of Claims 1-53 under 35 U.S.C. § 103(a)

Claims 1, 7-9, 11-13, 19-21, 29-33, 39-41, 45-47, 50 and 53 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Alfieri et al. (U.S. Patent No. 5,666,486) and Lennie et al. (U.S. Patent No. 6,092,213) in view of DeKoning et al. (U.S. Patent No. 6,073,218). These rejections are traversed.

Particular embodiments relate to a technique for maintaining a cluster definition for a network cluster, which can be accessed by member nodes without requiring network connectivity. The network cluster has at least two member nodes coupled to a shared repository.

¹See, e.g. Application, at pg. 4, ll. 13 (discussing that “[t]he present system includes a method and apparatus for providing common storage for the cluster definition which does not require network connectivity to propagate changes or to access the definition upon joining.”); *see also*, Application, at pg. 12, ll.25-29 (discussing that “[h]ere, however, only access to the shareable storage device is required to access the definition rather than network connectivity with the cluster.”).

A cluster definition for the network cluster is stored in the shared repository. A member node is a member in cluster and can access the cluster definition in the shared repository regardless of network connectivity. A member node proposes changes to the cluster definition by sending the changes to the shared repository. A coordinator node is selected from a member node to update the cluster definition at the shared repository to reflect the proposed change.

Although network connectivity may have been lost in the cluster, the nodes can access the cluster definition and propose changes to the cluster definition through disk-based messaging by proposing changes at the shared storage. As long as a node has access to the shareable storage, it can still participate in the network cluster. By way of comparison, prior systems required a member node to have network connectivity.

Alfieri and Lennie Require Network Connectivity

Like the cluster prior art, Alfieri and Lennie relate to cluster management systems that rely on network connectivity. In Alfieri, for example, if a member node loses network connectivity, its membership is revoked. In particular, the node is “forced” to leave the cluster. Likewise in Lennie, a member node is required to have access to the network in order to participate in the cluster and access the cluster definition. Furthermore, like the cluster prior art described in the application,² in order to determine the cluster definition, a member node in a network cluster of the Alfieri³ and Lennie⁴ type is required to have network connectivity with the

² “To determine the cluster definition, [in traditional cluster management systems,] a node would be required to have network connectivity with the cluster and would then be provided, by the network connection, the cluster definition for the network cluster. A drawback to this approach is that a node needs to have network connectivity with the cluster before the node can be provided with the definition for the cluster” *See* Application, pgs. 3-4. (Emphasis added.)

³ *See, e.g.* Alfieri, col. 7, ll. 56- col. 7, ll. 2 “An example of a forced leave is illustrated with respect to FIGS. 13-15 which begins with node N1 being joined to the cluster . . . If, by way of example, the membership managers of nodes N0 and N1 . . . have not received an acknowledge from the SFS subsystems 52 of their completion of node N0's graceful join processing because the SFS subsystems of nodes N0 and N1 cannot communicate due to an interconnect failure on node N0, node N0's graceful join is noted as in a hung state (FIG. 13). The membership manager of node N1 notices that node N1 can no longer communicate with node N0. Node N1 forces node N0 out of the cluster by marking node N0's state as forced-leaving. . . .”

⁴ *See, e.g.* Lennie, col. 2, ll. 44- col. 3, ll. 7.

cluster because it must be able to receive updates and propose changes to the cluster definition, which are distributed over the network connection to the member nodes.

With the present system, however, if a node's network connection fails, a node can still access the cluster definition and propose updates at the shared storage. As long as a node has access to the shared storage, the node can use disk-based messaging via the shared storage to communicate and participate in the cluster. In addition, regardless of network connectivity, a node in the claimed invention can propose changes to the cluster definition using disk-based messaging by sending its proposed changes to the shared repository. With this technique, the network cluster is not dependent on network connectivity. In this way, the inventive system provides a failsafe network cluster that can enable access to the most current update-to-date cluster definition and propose updates using the shared repository. Neither Alfieri nor Lennie discuss the claimed inventive concept, which is not dependent on network connectivity.

Furthermore, by relying on network connectivity and discussing that a node's membership is revoked if it loses network connectivity, Alfieri and Lennie teach away from the claimed invention, which is specifically directed to an approach that enables nodes to be members in the cluster regardless of whether they have network connectivity.

DeKoning Does Not Relate to a Network Cluster Where Nodes Are Joining and Exiting

The Examiner correctly acknowledges that Alfieri and Lennie do not teach the claimed concept of *accessing by a member nodes the cluster definition on the shared repository, regardless of network connectivity*. The Examiner cites DeKoning to show this claimed concept. DeKoning, however, relates to a RAID system that facilitates data redundancy by segmenting portions of stored data across several data disks. DeKoning provides RAID controllers that manage read/write access to storage nodes in the RAID system. Although DeKoning discusses the notion of a shared storage, DeKoning is not directed to a network cluster that is membership based. Other than the use of the term "node," DeKoning is not relevant to the claimed invention because DeKoning's "node" is not the same as the claimed "node."

In particular, DeKoning does not relate to a system that has different categories of nodes, potential members and members, e.g. nodes joining and exiting the cluster. DeKoning discusses storage nodes in the RAID system, but does not distinguish them into members and non-members.⁵ Conversely, the claimed invention is directed to a network cluster where nodes are joining and leaving. Independent Claims 33, 41, 46 and 50, for example, specifically require nodes that are *member nodes* and *potential member nodes*. Members and potential members can access the cluster definition regardless of network connectivity. Similarly, independent Claims 54, 62 and 67 require non-member and member nodes. This notion of members, potential or non-members does not exist in DeKoning because it is *not* directed to a network cluster in which a node's membership status is dynamic and, consequently, it does not relate the claimed network cluster.

DeKoning Does Not Relate to Avoiding Partitioned Network Clusters

Although the section in DeKoning cited by the Examiner discusses an approach for resolving communication failures, DeKoning does not address the types of failures that are specific to network clusters that are membership based. In particular, DeKoning is not directed to safeguards for avoiding partitioned network clusters. Instead, DeKoning relates to preventing data loss by providing disk controller to coordinate access to the shared storage to facilitate synchronization of data across the RAID drives.

By way of comparison, the invention relates to safeguards for preventing partitioned network clusters. As discussed in the application and cluster membership prior art, such as Alfieri and Lennie, cluster partitioning results when the member nodes begin operating using an invalid cluster definition and this causes the cluster to degenerate into multiple cluster partitions including a subset of the cluster network nodes, each cluster partition operating independently of each other. These partitions may be the result of miscommunication that occurs from nodes joining or exiting the network cluster, the so-called partition-in-time problem. One important safeguard provided by the invention via the shared repository is enabling member nodes to use

⁵See e.g. DeKoning, Col. 11, ll. 24-41 (discussing the communications between nodes in the disk array, without distinguishing between different types of nodes).

disk-based messaging in that a shared repository enables the members and potential members to communicate and access the cluster definition on the shared repository. By using disk-based messaging via the shared repository, nodes are able to access the cluster definition even if they have lost network connectivity.

DeKoning does not attempt to address the partition-in-time problem in such membership-based network clusters because DeKoning does not relate to managing the cluster definition in clusters that are vulnerable to the possibility of miscommunication resulting from nodes joining or exiting. Thus, DeKoning does not relate to clusters that are membership based and, consequently, DeKoning is nonanalogous art. In this way, there is no motivation to combine DeKoning with Alfieri and Lennie because Alfieri and Lennie relate to membership-based clusters.

The Combination of Alfieri, Lennie and DeKoning Does Not Result in the Claimed Invention

Nevertheless, even if DeKoning is analogous art and there is motivation to combine the references, the combination of DeKoning with Alfieri and Lennie would not result in the claimed invention. In particular, the combination would result in an in a network cluster that is still dependent on network connectivity. Alfieri and Lennie discuss approaches that rely on the network interconnects to communicate the cluster definition to the member nodes, and they base membership on network connectivity; while DeKoning does not relate to membership at all, but rather relates to a RAID system that uses RAID controllers to coordinate access to its shared storage drives. Thus, the resulting combined system would involve a network cluster that distributes cluster definitions over the network to its members who also share access to a RAID shared storage. There is nothing about this combined system that would eliminate the dependency on the network infrastratucutre for member nodes to access the cluster definition, to propose changes to the cluster defintion and to receive updates to the cluster definition.

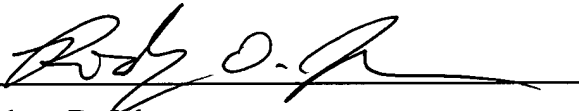
Therefore, Alfieri, Lennie and DeKoning, taken separately, or in combination, do not suggest the requirements of independent Claims 1, 11, 13, 32, 33, 41, 46, 50, 53, 54, 62 and 67. Thus, the rejections of independent Claims 1, 11, 13, 32, 33, 41, 46, 50, 53, 54, 62 and 67 under § 103(a) and their respective dependent claims should be withdrawn.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned attorney.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By 

Rodney D. Johnson
Registration No. 36,558
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133

Dated: *March 1, 2005*